UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/783,865	02/20/2004	Greg Pellegrino	200301878-2	7590

7590 02/09/2009

HEWLETT-PACKARD COMPANY
Intellectual Property Administration
P.O. Box 272400
Fort Collins, CO 80527-2400

EXAMINER	
CHU, GABRIEL L	
ART UNIT	PAPER NUMBER
2114	

MAIL DATE	DELIVERY MODE
02/09/2009	PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Office Action Summary	Application No.	Applicant(s)	
	10/783,865	PELLEGRINO ET AL.	
	Examiner	Art Unit	
	Gabriel L. Chu	2114	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE *3* MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1)☒ Responsive to communication(s) filed on *08 December 2008*.

2a)☐ This action is **FINAL**. 2b)☒ This action is non-final.

3)☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

4)☒ Claim(s) *21-41* is/are pending in the application.

4a) Of the above claim(s) _____ is/are withdrawn from consideration.

5)☐ Claim(s) _____ is/are allowed.

6)☒ Claim(s) *21-41* is/are rejected.

7)☐ Claim(s) _____ is/are objected to.

8)☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

9)☐ The specification is objected to by the Examiner.

10)☒ The drawing(s) filed on *20 February 2004* is/are: a)☒ accepted or b)☐ objected to by the Examiner.

Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

11)☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

12)☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

a)☐ All b)☐ Some * c)☐ None of:

1.☐ Certified copies of the priority documents have been received.

2.☐ Certified copies of the priority documents have been received in Application No. _____.

3.☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) ☐ Notice of References Cited (PTO-892)
2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
3) ☐ Information Disclosure Statement(s) (PTO/SB/08) Paper No(s)/Mail Date _____.
4) ☐ Interview Summary (PTO-413) Paper No(s)/Mail Date. _____ .
5) ☐ Notice of Informal Patent Application
6) ☐ Other: _____.

DETAILED ACTION

Claim Rejections - 35 USC § 112

1. The following is a quotation of the second paragraph of 35 U.S.C. 112:

> The specification shall conclude with one or more claims particularly pointing out and distinctly claiming the subject matter which the applicant regards as his invention.

2. **Claims 21-37 rejected under 35 U.S.C. 112, second paragraph, as being indefinite for failing to particularly point out and distinctly claim the subject matter which applicant regards as the invention.**

3. Referring to claim 21, and subsequently claims 22-25, Applicant has claimed “an available direct communication link between the first compute node, a first switch, and the first storage node”. It is not clear how a single link can “directly” connect all three components, and it is further unclear how the intended path between a first compute to first switch to first storage forms a "direct" path by a common interpretation of the term. This is understood to refer to a “an available direct communication link from the first compute node through a first switch to the first storage node” as similarly claimed in the last limitation describing the “direct” path from the second compute node through the first/second switch to the first storage node.

While Examiner has found no specific definition, Examiner supposes that from claim 21, for example, in further view of paragraph 12 of Applicant's pre-grant publication, that Applicant intends that a switch, or at least something like a switch, is not considered an element that will make a path "indirect", whereas a path taken via, for example, a “compute node” would be an indirect path. Examiner proposes the above

amendment for at least some consistency and in view of Applicant's description (for example, claim 21) of a path via a switch as "direct".

4. Referring to claim 26, claims 27-30, similar to claim 21 above, the claim uses the term "direct" in manner that is unclear.

5. Referring to claim 26, claims 27-30, in the second to last line, "second node" is understood to refer to "second compute node".

6. Referring to claim 30, "second node" is understood to refer to "second compute node".

7. Referring to claim 31, claims 32-34, “the storage network” is understood to refer to “the storage area network”.

8. Referring to claims 35-37, claim 35 uses “direct" as well, and is similarly unclear as noted above for claim 21.

Claim Rejections - 35 USC § 103

9. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

> (a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negatived by the manner in which the invention was made.

10. Claims 21-24, 35-37 rejected under 35 U.S.C. 102(b) as being anticipated by US 5513314 to Kandasamy et al. in view of Official Notice.

11. Referring to claim 21 Kandasamy discloses a method of managing a write request from a first compute node in a storage network to a first storage node in the storage network, comprising:

if there is an available "direct" communication path between the first compute node and the first storage node, then executing the write request from the first compute node to the first storage node using the available direct communication link (From line 20 of column 6, "The preferred approach to enabling proxy operation is to establish a virtual server system formed as a logical composite of a primary and one or more secondary servers, constituting an active group, relative to a specific set of mutually fault tolerantly protected individual filesystems. The virtual server is identified by a unique hostname and IP address. A MAC layer multicast address, commonly subscribed to and accessible by the primary and secondary servers, is also assigned to the virtual server. Client workstations access the fault tolerant protected filesystems of the active group by reference to the IP address and through the MAC multicast address of the virtual server. Consequently, NFS requests directed to the virtual server are commonly received by the primary and secondary servers of the active group." Further, see figure 1, where clients and servers are connected via a LAN 16, "directly".);

if there is not an available communication path between the first compute node and the first storage node (In Kandasamy, these would be the client and the secondary server), then:

transmitting the write request from the first compute node to a second compute node (In Kandasamy, the first server) if there is an available "direct" communication

path from the first compute node to the second compute node (In Kandasamy, the path from the client to the primary server) and an available "direct" communication link from the second compute node to the first storage node (From line 65 of column 12 (with emphasis), "Alternately, the primary server 12, prior to issuing a completion datagram for the NFS write request, may itself initiate a conventional NFS write operation **directly** to the secondary server 14 to force completion of the client requested NFS write operation." Wherein it is "the" write request issued by the client.).

Although Kandasamy does not specifically disclose that such a storage network may be a SAN, SANs are very well known in the art. Examiner takes official notice for SANs. A person having ordinary skill in the art at the time of the invention could have been motivated to use a SAN because it allows storage devices appears as if they are locally attached to the operating system, and furthermore, the implementation of a SAN is not technology specific and may include such technologies as ATA, SCSI, Ethernet, and fibre channel.

Further, although Kandasamy does not specifically disclose that a switch may connect nodes (for example between a first compute node and a first storage node or between a second compute node and a first storage node), switches used to connect nodes in a network are very well known in the art. Examiner takes official notice for a switch. A person having ordinary skill in the art at the time of the invention could have been motivated to use a switch because it allows network segments to be connected.

12. Referring to claim 22, Kandasamy discloses if executing the write request from the first compute node to the first storage node generates a timeout failure, then:

transmitting the write request from the first compute node to a second compute node if there is an available communication path from the first compute node to the second compute node and an available communication path from the second compute node to the first storage node (From line 19 of column 12, “The primary server 12 may fail to receive an acknowledge datagram from the secondary server 14 for a number of reasons. The secondary server 14, in a first instance, may have failed to properly receive the client write request. Alternately, the secondary server 14, in performing the request, may fail for some reason to complete the request. In either case, no acknowledgment datagram is issued by the secondary server 14. Another possibility is that the client write request was properly received and performed by the secondary server 14, but the issued acknowledgment datagram was not properly received by the primary server 12. In each of these events, the primary server 12 is left sleeping on the DRC entry for an acknowledgment datagram that is not received. However, in accordance with the present-invention, a sleep timer is set by the primary server 12 in putting the nfsd process to sleep on DRC entry. The nfsd process awakes 86 on timeout of the sleep timer in the absence of any received acknowledge datagram. Alternately, the sleep timer is effectively expired upon the aging of the DRC entry through operation of the DRC-LRU algorithm. In either event, the primary server 12 then transitions to a backup failure recovery mode 88.”).

13. Referring to claim 23, Kandasamy discloses transmitting the write request from the first compute node to the second compute node comprises encapsulating the write request (From line 31 of column 7, “On both the primary files server 12 and secondary

file server 14, the datagram representing the NFS write request is processed by a substantially conventional TCP/IP stack. In relevant part, this network stack includes a physical layer, a data link layer, a network layer, a transport layer, a session layer and an. application layer.").

14. Referring to claim 24, Kandasamy discloses executing the write request from the second compute node to the first storage node (From line 65 of column 12, "Alternately, the primary server 12, prior to issuing a completion datagram for the NFS write request, may itself initiate a conventional NFS write operation directly to the secondary server 14 to force completion of the client requested NFS write operation.").

15. Referring to claim 35, see rejection of claim 21.

16. Referring to claim 36, see rejection of claim 22.

17. Referring to claim 37, see rejection of claim 23.

18. Claim 25 rejected under 35 U.S.C. 103(a) as being unpatentable over US 5513314 to Kandasamy et al. and Official Notice as applied to claim 24 above, and further in view of Official Notice.

19. Referring to claim 25, although Kandasamy does not specifically disclose transmitting an error message from the second compute node to the first compute node if the write request fails, indicating such a failure is very well known in the art. Examiner takes official notice for sending a failure indication. A person of ordinary skill in the art at the time of the invention could have been motivated to send such an indication because, as shown in line 59 of column 12, "the primary server 12 may intentionally withhold issuing any completion datagram to the client 26", and such an indication

would let the client know why it has not yet received a completion datagram. Further, from line 31 of column 13, "Once a failover event is declared, the primary server 12 ultimately issues a completion datagram to the client workstation 26 indicating a proper completion of the pending NFS write request. The primary server 12 also preferably logs the failover event and issues appropriate status messages to the administrator's console." In either case, such an indication would either inform the client of a failover or, if the administrator is the client, then clearly this would be done anyway.

20. Claim 26-30 rejected under 35 U.S.C. 103(a) as being unpatentable over US 5513314 to Kandasamy et al. in view of US 6751190 to Swallow and Official Notice.

21. Referring to claim 26, Kandasamy discloses a method of managing a write request from a first compute node in a storage network to a mirrored storage data set having a first storage node and a second storage node in the storage network, comprising:

if there are available communication paths between the first compute node and both the first storage node and the second storage node in the mirrored data set, then executing the write request from the first compute node to both the first storage node and the second storage node using the available communication paths (From line 20 of column 6, "The preferred approach to enabling proxy operation is to establish a virtual server system formed as a logical composite of a primary and one or more secondary servers, constituting an active group, relative to a specific set of mutually fault tolerantly protected individual filesystems. The virtual server is identified by a unique hostname

and IP address. A MAC layer multicast address, commonly subscribed to and accessible by the primary and secondary servers, is also assigned to the virtual server. Client workstations access the fault tolerant protected filesystems of the active group by reference to the IP address and through the MAC multicast address of the virtual server. Consequently, NFS requests directed to the virtual server are commonly received by the primary and secondary servers of the active group.");

if there are no available communication paths between the first compute node and the first storage node and the second storage node, then invoking an error routine (From line 19 of column 12, "The primary server 12 may fail to receive an acknowledge datagram from the secondary server 14 for a number of reasons. The secondary server 14, in a first instance, may have failed to properly receive the client write request. Alternately, the secondary server 14, in performing the request, may fail for some reason to complete the request. In either case, no acknowledgment datagram is issued by the secondary server 14. Another possibility is that the client write request was properly received and performed by the secondary server 14, but the issued acknowledgment datagram was not properly received by the primary server 12. In each of these events, the primary server 12 is left sleeping on the DRC entry for an acknowledgment datagram that is not received. However, in accordance with the present-invention, a sleep timer is set by the primary server 12 in putting the nfsd process to sleep on DRC entry. The nfsd process awakes 86 on timeout of the sleep timer in the absence of any received acknowledge datagram. Alternately, the sleep timer is effectively expired upon the aging of the DRC entry through operation of the

DRC-LRU algorithm. In either event, the primary server 12 then transitions to a backup failure recovery mode 88.");

if there is an available communication path between the first compute node and only one of the first storage node and the second storage node in the mirrored data set (See figures 4, 5 of Kandasamy wherein it is disclosed that either one of the paths may fail.), then: executing the write request from the first compute node to the first storage node or the second storage node via the available communication path (Wherein whichever path that is not failed results in a successful write.).

Kandasamy further discloses that a surrogate write may occur through a proxy node (From line 65 of column 12, "Alternately, the primary server 12, prior to issuing a completion datagram for the NFS write request, may itself initiate a conventional NFS write operation directly to the secondary server 14 to force completion of the client requested NFS write operation." Wherein it is "the" write request issued by the client.).

Although Kandasamy does not specifically disclose transmitting the write request from the first compute node to a second compute node if there is an available "direct" communication path from the first compute node to the second compute node and an available communication path from the second compute node through a first switch or a second switch to the first storage node or the second storage node, routing around failure via an intermediary node (Here, taken to be the "second compute node", for example Figure 1, node D.) which may further use a switch (Here, taken to be the "first switch" or "second switch", for example in figure 1, on a path 102 to 110, a node such as C which lies between D and 110.) is very well known in the field of networking. An

example of this is shown by Swallow, from line 1 of column 8, "FIG. 9 shows the steps performed by intermediate node_A 104 to reroute data packets for the primary tunnel 126 along the bypass tunnel 128. Referring to FIG. 9, in step 900 intermediate node_A 104 discovers that it can not forward data packets to intermediate node B 106 (FIG. 10) because of a communication link failure. In step 902 intermediate node_A 104 (FIG. 1) determines from its saved copy of the Record Route Object 416 (FIG. 6G) in the label table that intermediate node_C 108 (FIG. 1) is adjacent to intermediate node_B 106 (FIG. 1) in the primary tunnel 126 (FIG. 1). Intermediate node_A 104 (FIG. 1) also determines from the label table the incoming label 802 (FIG. 8A) assigned by intermediate node_C 108 (FIG. 1) to be transmitted with data packets from intermediate node_B 106 (FIG. 1) to intermediate node_C 108 (FIG. 1). In step 904 intermediate node_A 104 (FIG. 1) establishes a bypass tunnel 128 through intermediate node_D 120 to intermediate node_C 108 using the same method for establishing the primary tunnel 126 described in conjunction with FIGS. 2 and 3." A person of ordinary skill in the art at the time of the invention could have been motivated to use such a bypass tunnel because, as disclosed in both Kandasamy and Swallow, there is a need for routing around failure in a network. From line 26 of column 14 of Kandasamy, "Where a secondary server 14 looses all communication with an active primary server 12, the secondary server 14 should not assume that the primary server 12 has failed. Rather, the likely failure point is the LAN 16. In that case, the possibility of multiple active primary servers 12 serving the same file system must be avoided. Therefore, a secondary server 14 should not be promoted to an active primary state where all

communication with an active primary server 12 is lost." From the abstract of Swallow, "A bypass tunnel to bypass a node in the pre-defined path may be established to reroute data packets around a failed communication link in the tunnel. The bypass tunnel may be established before the communication link failure providing fast tunnel restoration."

Although Kandasamy and Swallow do not specifically disclose that such a storage network may be a SAN, SANs are very well known in the art. Examiner takes official notice for SANs. A person having ordinary skill in the art at the time of the invention could have been motivated to use a SAN because it allows storage devices appears as if they are locally attached to the operating system, and furthermore, the implementation of a SAN is not technology specific and may include such technologies as ATA, SCSI, Ethernet, and fibre channel.

22. Referring to claim 27, Kandasamy in view of Swallow discloses if executing the write request from the first compute node to the first storage node generates a **timeout** failure, then: transmitting the write request from the first compute node to a second compute node if there is an available communication path from the first compute node to the second compute node and an available communication path from the second source node to the first storage node (From line 19 of column 12 of Kandasamy, "The primary server 12 may fail to receive an acknowledge datagram from the secondary server 14 for a number of reasons. The secondary server 14, in a first instance, may have failed to properly receive the client write request. Alternately, the secondary server 14, in performing the request, may fail for some reason to complete the request. In

either case, no acknowledgment datagram is issued by the secondary server 14. Another possibility is that the client write request was properly received and performed by the secondary server 14, but the issued acknowledgment datagram was not properly received by the primary server 12. In each of these events, the primary server 12 is left sleeping on the DRC entry for an acknowledgment datagram that is not received. However, in accordance with the present-invention, a sleep timer is set by the primary server 12 in putting the nfsd process to sleep on DRC entry. The nfsd process awakes 86 on timeout of the sleep timer in the absence of any received acknowledge datagram. Alternately, the sleep timer is effectively expired upon the aging of the DRC entry through operation of the DRC-LRU algorithm. In either event, the primary server 12 then transitions to a backup failure recovery mode 88." Further, see figures 4, 5 of Kandasamy, wherein either path may fail.).

23. Referring to claim 28, Kandasamy in view of Swallow discloses executing the write request from the second compute node to the first storage node (See bypass tunnel.).

24. Referring to claim 29, see rejection of claim 27.

25. Referring to claim 30, see rejection of claim 28.

26. Claims 31-33, 38-40 rejected under 35 U.S.C. 103(a) as being unpatentable over US 6751190 to Swallow in view of US 5513314 to Kandasamy et al. and Official Notice.

27. Referring to claim 31, Swallow discloses receiving, at a second compute node, a query from a first compute node, wherein the query identifies a target node in the

storage network; transmitting a reply to the first compute node, wherein the reply includes a signal component indicating there is an available communication path between the second compute node and the target node (See figures 2-4. Further, from line 19 of column 4, “Referring to FIG. 4, the Message Identification Object 402 identifies the message type as a Path Message 400. The Session Object 404 identifies the session. The format of the Session Object 404 is shown in FIG. 6A. Referring to FIG. 6A, the Session Object 404 provides the IPv4 Endpoint Address 602 for the receive endpoint 110 (FIG. 1), a Tunnel Session Identification 604 indicating a value for the primary tunnel 126 (FIG. 1) and an Extended Tunnel Identification 630 for further defining the session. Continuing with FIG. 4, the RSVP_HOP Object 406 includes the address of the transmit endpoint 102 (FIG. 1) in IPv4 format. The Explicit Route Object 408 includes a list of the IPv4 addresses for the intermediate nodes 104, 106, 108 in the primary tunnel 126 (FIG. 1). The format of the Explicit Route Object 408 is shown in FIG. 6D. Referring to FIG. 6D, to set up an explicit route between the transmit endpoint 102 (FIG. 1) and the receive endpoint 110 (FIG. 1), the IPv4 address for intermediate node_A 104 (FIG. 1) is stored in IPv4 address 610, the IPv4 address for intermediate node_B 106 (FIG. 1) is stored in IPv4 address 612, and the IPv4 address for intermediate node_C 108 (FIG. 1) is stored in IPv4 address 614. Continuing with FIG. 4, the Label Request Object 410 indicates that the transmit endpoint 110 (FIG. 1) is requesting a label to be assigned for communication link 112 (FIG. 1) between the transmit endpoint 110 (FIG. 1) and intermediate node_A 104. The transmit endpoint 110 (FIG. 1) transmits the returned assigned label with each data packet associated with the

application data to be transmitted on the primary tunnel 126 (FIG. 1). The format of the Label Request Object 410 is shown in FIG. 6E. ").

Although Swallow does not specifically disclose that such a (bypass) tunnel may be used for the surrogate write operation and relaying write operations from the first compute node to the target node, such a use is known in the art. An example of this is shown by Kandasamy. From line 20 of column 6, "The preferred approach to enabling proxy operation is to establish a virtual server system formed as a logical composite of a primary and one or more secondary servers, constituting an active group, relative to a specific set of mutually fault tolerantly protected individual filesystems. The virtual server is identified by a unique hostname and IP address. A MAC layer multicast address, commonly subscribed to and accessible by the primary and secondary servers, is also assigned to the virtual server. Client workstations access the fault tolerant protected filesystems of the active group by reference to the IP address and through the MAC multicast address of the virtual server. Consequently, NFS requests directed to the virtual server are commonly received by the primary and secondary servers of the active group." From line 65 of column 12, "Alternately, the primary server 12, prior to issuing a completion datagram for the NFS write request, may itself initiate a conventional NFS write operation directly to the secondary server 14 to force completion of the client requested NFS write operation." A person of ordinary skill in the art at the time of the invention could have been motivated to perform such a surrogate write operation because, as disclosed in Swallow it is known that such a network failure may impede normal network function, and as disclosed by Kandasamy, one such

functionality may be a write.

Although Kandasamy and Swallow do not specifically disclose that such a storage network may be a SAN, SANs are very well known in the art. Examiner takes official notice for SANs. A person having ordinary skill in the art at the time of the invention could have been motivated to use a SAN because it allows storage devices appears as if they are locally attached to the operating system, and furthermore, the implementation of a SAN is not technology specific and may include such technologies as ATA, SCSI, Ethernet, and fibre channel.

28. Referring to claim 32, Swallow in view of Kandasamy discloses determining whether there is an available communication path between the second compute node and the target node (See figures 2-4. Further, from line 19 of column 4, "Referring to FIG. 4, the Message Identification Object 402 identifies the message type as a Path Message 400. The Session Object 404 identifies the session. The format of the Session Object 404 is shown in FIG. 6A. Referring to FIG. 6A, the Session Object 404 provides the IPv4 Endpoint Address 602 for the receive endpoint 110 (FIG. 1), a Tunnel Session Identification 604 indicating a value for the primary tunnel 126 (FIG. 1) and an Extended Tunnel Identification 630 for further defining the session. Continuing with FIG. 4, the RSVP_HOP Object 406 includes the address of the transmit endpoint 102 (FIG. 1) in IPv4 format. The Explicit Route Object 408 includes a list of the IPv4 addresses for the intermediate nodes 104, 106, 108 in the primary tunnel 126 (FIG. 1). The format of the Explicit Route Object 408 is shown in FIG. 6D. Referring to FIG. 6D, to set up an explicit route between the transmit endpoint 102 (FIG. 1) and the receive endpoint 110 (FIG. 1),

the IPv4 address for intermediate node_A 104 (FIG. 1) is stored in IPv4 address 610, the IPv4 address for intermediate node_B 106 (FIG. 1) is stored in IPv4 address 612, and the IPv4 address for intermediate node_C 108 (FIG. 1) is stored in IPv4 address 614. Continuing with FIG. 4, the Label Request Object 410 indicates that the transmit endpoint 110 (FIG. 1) is requesting a label to be assigned for communication link 112 (FIG. 1) between the transmit endpoint 110 (FIG. 1) and intermediate node_A 104. The transmit endpoint 110 (FIG. 1) transmits the returned assigned label with each data packet associated with the application data to be transmitted on the primary tunnel 126 (FIG. 1). The format of the Label Request Object 410 is shown in FIG. 6E. “).

29. Referring to claim 33, Swallow in view of Kandasamy discloses relaying write operations from the compute node to the target node comprises: receiving an encapsulated write request from the first compute node; de-encapsulating the encapsulated write request; and executing the write request from the second node to the target node (From line 31 of column 7 of Kandasamy, “On both the primary files server 12 and secondary file server 14, the datagram representing the NFS write request is processed by a substantially conventional TCP/IP stack. In relevant part, this network stack includes a physical layer, a data link layer, a network layer, a transport layer, a session layer and an. application layer.” From Swallow, figures 7, 9, line 30 of column 8, “In step 908 intermediate node_A 104 redirects data packets for the primary tunnel 126 (FIG. 1) through the bypass tunnel 128 (FIG. 1) to intermediate node_C 108 (FIG. 1). Intermediate node_A 104 uses the incoming label 802 (FIG. 8A) from transmitting node 102 (FIG. 1) to index into intermediate node_A's 104 incoming label

table 804 (FIG. 1). Intermediate node_A 104 adds the incoming label 802 (FIG. 8A) for intermediate node_C 108 for the primary tunnel 126 to the label stack associated with the data packet to be forwarded on the primary tunnel 126. Intermediate node_A 104 (FIG. 1) adds the outgoing label value 820 (FIG. 8B) from the NHLFE 808 in the incoming label table 804 (FIG. 1) to the top of the label stack and forwards the data packet and the associated label stack to intermediate node_D 120. Returning to FIG. 7, in step 702 intermediate node_D 120 receives the data packet and associated label stack from intermediate node_A 104 and uses the incoming label value 806 on the top of the label stack to determine the outgoing label value 820 to be placed on the label stack. The incoming label value 806 is an index to a NHLFE 808 in the incoming label map 804 stored in intermediate node_D 120. Intermediate node_D determines from the label stack operation field 816 in the NHLFE 808 in the incoming label map 804 to pop the top label off the label stack. In step 706 intermediate node_D 120 forwards the modified label stack and the data packet to intermediate node_C 108. In step 702 intermediate node_C 108 (FIG. 1) uses the incoming label value 802 at the top of the label stack associated with the data packet forwarded from intermediate node_D 108 (FIG. 1). The incoming label value 806 is the same as the incoming label value 806 forwarded from intermediate node_B 106 (FIG. 1) for the primary tunnel 126. In step 704 intermediate node_C 108 (FIG. 1) pops the incoming label 802 from the label stack according to the label stack operation field 816 in the NHLFE 808 for the incoming label 802 from intermediate node_D 120 (FIG. 1). Intermediate node_C 108 (FIG. 1) seeing the same incoming label value 806 performs the same operation on the data packet's

label stack that it performs for the data packet's label stack if the data packet was transferred along the primary tunnel 126 (FIG. 1). In step 706 intermediate node_C 108 (FIG. 1) forwards the data packet with no associated label stack to the receive endpoint 110 (FIG. 1) of the primary tunnel 126 (FIG. 1).".).

30. Referring to claim 38, see rejection of claim 31.

31. Referring to claim 39, see rejection of claim 32.

32. Referring to claim 40, see rejection of claim 33.

33. Claim 34, 41 rejected under 35 U.S.C. 103(a) as being unpatentable over US 6751190 to Swallow and US 5513314 to Kandasamy et al. and Official Notice as applied to claim 31, 38 above, and further in view of Official Notice.

34. Referring to claim 34, although Swallow in view of Kandasamy does not specifically disclose transmitting failure signal from the second source node to the first source node if the write request fails, indicating such a failure is very well known in the art. Examiner takes official notice for sending a failure indication. A person of ordinary skill in the art at the time of the invention could have been motivated to send such an indication because, as shown in line 59 of column 12 of Kandasamy, "the primary server 12 may intentionally withhold issuing any completion datagram to the client 26", and such an indication would let the client know why it has not yet received a completion datagram. Further, from line 31 of column 13, "Once a failover event is declared, the primary server 12 ultimately issues a completion datagram to the client workstation 26 indicating a proper completion of the pending NFS write request. The primary server 12 also preferably logs the failover event and issues appropriate status messages to the

administrator's console." In either case, such an indication would either inform the client of a failover or, if the administrator is the client, then clearly this would be done anyway.

35. Referring to claim 41, see rejection of claim 34.

Double Patenting

36. The nonstatutory double patenting rejection is based on a judicially created doctrine grounded in public policy (a policy reflected in the statute) so as to prevent the unjustified or improper timewise extension of the "right to exclude" granted by a patent and to prevent possible harassment by multiple assignees. A nonstatutory obviousness-type double patenting rejection is appropriate where the conflicting claims are not identical, but at least one examined application claim is not patentably distinct from the reference claim(s) because the examined application claim is either anticipated by, or would have been obvious over, the reference claim(s). See, e.g., *In re Berg*, 140 F.3d 1428, 46 USPQ2d 1226 (Fed. Cir. 1998); *In re Goodman*, 11 F.3d 1046, 29 USPQ2d 2010 (Fed. Cir. 1993); *In re Longi*, 759 F.2d 887, 225 USPQ 645 (Fed. Cir. 1985); *In re Van Ornum*, 686 F.2d 937, 214 USPQ 761 (CCPA 1982); *In re Vogel*, 422 F.2d 438, 164 USPQ 619 (CCPA 1970); and *In re Thorington*, 418 F.2d 528, 163 USPQ 644 (CCPA 1969).

A timely filed terminal disclaimer in compliance with 37 CFR 1.321(c) or 1.321(d) may be used to overcome an actual or provisional rejection based on a nonstatutory double patenting ground provided the conflicting application or patent either is shown to be commonly owned with this application, or claims an invention made as a result of activities undertaken within the scope of a joint research agreement.

Effective January 1, 1994, a registered attorney or agent of record may sign a terminal disclaimer. A terminal disclaimer signed by the assignee must fully comply with 37 CFR 3.73(b).

37. **Claims 21-24, 26-33, 35-40 rejected on the ground of nonstatutory obviousness-type double patenting as being unpatentable over claims 1-15 of U.S. Patent No. 6725393.** Although the conflicting claims are not identical, they are not patentably distinct from each other. Claims 21-24, 26-33, 35-40 of the instant application are anticipated by claims 1-15 of U.S. Patent No. 6725393 in that claims 1-15 of U.S. Patent No. 6725393 contain all of the limitations of claims 21-24, 26-33, 35-40 of the instant application. Claims 21-24, 26-33, 35-40 of the instant application therefore are not patently distinct from the earlier patent claims, and as such are

unpatentable for obvious-type double patenting. (In re Goodman (CAFC) 29 USPQ2d 2010). While limitations of the claims of U.S. Patent No. 6725393 may be broader than the claims of the instant application, the language and the disclosure of U.S. Patent No. 6725393 indicate that the limitation of claims of the instant application are merely a subset of U.S. Patent No. 6725393. These differences are not sufficient to render the claims patentably distinct. Georgia-Pacific Corp. v. United States Gympsum Co., 195 F.3d 1322, 1325, 52 USPQ2d 1590, 1593 (Fed. Cir. 1999).

38. **Claim 25, 34, 41 rejected on the ground of nonstatutory obviousness-type double patenting as being unpatentable over claims 1-15 of U.S. Patent No. 6725393 in view of Official Notice.** See reasoning above. Further, although U.S. Patent No. 6725393 does not specifically disclose transmitting failure signal from the second source node to the first source node if the write request fails, indicating such a failure is very well known in the art. Examiner takes official notice for sending a failure indication. A person of ordinary skill in the art at the time of the invention could have been motivated to send such an indication because a failure has occurred, and signaling such would have been informative for any number of reasons including logging, recovery, or just notification.

39. Note that double patenting is merely being applied because the terminal disclaimer has not yet been processed. After approval, the rejection will be removed.

Response to Arguments

40. Applicant's arguments filed 8 December 2008 have been fully considered but they are not persuasive. Applicant is largely relying on the use of the terms “direct” and

"compute" and specifying the storage network as a SAN. In some claims, *only* the SAN amendment and the term "compute" are used.

41. Firstly, that a storage network environment may be a SAN is not seen to materially differentiate over the prior art. SAN provides an additional layer of transparency intended to work with existing technologies. A person having ordinary skill in the art would have recognized that a generic storage network could have been implemented as a SAN. While it certainly would have taken effort, this is not inventive effort.

42. Secondly, the term "direct" has problems, as was discussed in an interview with Attorney. Examiner has hypothesized that Applicant has intended for network nodes, or some such, not to count when considering directness of a path, although even then, whether something is considered network hardware is not clear (see for example, paragraph 5 of the pre-grant publication). At any rate, as noted by the 112 rejections above, the use of the term "direct" is, at best, not clear.

43. Thirdly, the term "compute node", intended by Applicant to differentiate beyond merely a "source", is still a broad term. A compute node is, in view of the art, merely a node that has/can/for/etc... computing. As a node is already minimally a computer, applying the label of "compute" could actually be a broader term than a "source" node, which at least implied that it was the source of something. Looking to Applicant's specification, for example paragraph 5 of the pre-grant publication, it can be seen that "compute" nodes are supposed to be different from storage nodes and, perhaps, "general-purpose network functions". However, from at least the same paragraph, it is

clear that Applicant intended the term to be broad, as no specific functionality is assigned to it beyond somehow “utilizing separate network hardware” and “typically... use storage provided by storage nodes" and "may, and often do, have additional storage devices directly attached". Indeed, it is unclear what a compute node must be as it is minimally interpreted as a node on a network that may or may not use a storage node and may or may not have storage attached.

44. Note, Examiner is not arguing that these terms do not mean anything and that SAN cannot differentiate, but to the extent applied/claimed, they provide minimal differentiation. In terms of inventive concept, they do not appear to differentiate significantly beyond merely routing around a failed path in a type of network with types of nodes.

Conclusion

45. The prior art made of record and not relied upon is considered pertinent to applicant's disclosure. See notice of references cited.

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Gabriel L. Chu whose telephone number is (571) 272-3656. The examiner can normally be reached on weekdays between 8:30 AM and 5:00 PM.

If attempts to reach the examiner by telephone are unsuccessful, the examiner’s supervisor, Scott Baderman can be reached on (571) 272-3644. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see http://pair-direct.uspto.gov. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

/Gabriel L. Chu/
Primary Examiner
Art Unit 2114

gc